

2/PRTS

An access control mechanism for packet switched communication networks

1. Field of invention

The invention relates to media access control (MAC) over a shared multiplexed link using packet transmission. It is particularly directed to the management of the transfer of information from services which stipulate certain delays and delay variations, such as circuit emulated traffic, in a packet switched network, and specifically between a central node and one of several multipoint nodes over a shared medium.

2. Background art

Media access control (MAC) in a packet switched network is a mechanism for determining which of a plurality of nodes is allowed to send information over a shared medium. In a point-to-multipoint configuration it is common that the head end or central node in a shared link will control the utilisation of the link by multipoint nodes. This is achieved by polling the various multipoint nodes. The polling order and frequency is generally controlled according to traffic demand. In a time division multiplexed TDM link, time slots are allocated by the head end of a link in response to requests sent by the multipoint nodes when traffic is awaiting transmission at these node. The slots are allocated using permits that designate which time slot in the future traffic stream may be utilised for which service traffic. When more traffic is waiting than there are available time slots, the central node will issue queuing and prioritising requests to determine which service should be allocated channel capacity first, and which can wait.

Packet switched networks are increasingly utilised to transport a variety of traffic, some of which is traditionally circuit switched. An example of this is the transmission of voice over an internet protocol (IP) based network. While speech information may be packetised, certain transmission delays and delay variations must be respected if the speech service is to be acceptable to the end user. This is a general requirement of all circuit emulated traffic. Accordingly, with the above-described request based protocols, a multipoint node carrying circuit emulated traffic requires permits to be allocated without delay and with high priority. When more than one delay sensitive service utilises the shared medium the access control becomes more complex. This situation is made still more difficult when the delay sensitive services use different frame lengths and possibly also different bit rates so that intervals between the required high priority permits are always changing. The result is often excessive delays in one or more of the delay sensitive services leading to the end user experiencing an unacceptable quality of service.

SUMMARY OF INVENTION

It is therefore an object of the invention to provide a method and arrangement enabling the efficient handling of hybrid traffic over a shared multiplexed link.

It is a further object of the present invention to provide a method and arrangement enabling the handling of several delay sensitive services over a shared medium in a packet switched network.

These and further objects are achieved in a method and arrangement for controlling access to a central node over shared medium by one of several network nodes carrying both delay sensitive and non-delay sensitive traffic, by

providing a scheduler that defines a predetermined scheduling interval for the delay sensitive traffic services. The scheduling interval corresponding essentially to the time in terms of time slots for assembling a data packet of the delay sensitive traffic service. Prior to allocating a time slot, the central node consults the scheduler to determine if a scheduling interval has elapsed. If an interval has elapsed, a time slot is allocating to the network node carrying the corresponding delay sensitive traffic service. If no scheduling interval has elapsed for the time slot consulted, the central node allocates a time slot to a non-delay sensitive service on the basis of stored information concerning the amount of traffic awaiting transmission at the various network nodes.

The scheduler is a cyclical storage means that is divided into cells representing a time slot on the shared medium and adapted to contain scheduling information relating to the traffic service.

The scheduler may comprise scheduling information for several services in a single cyclical storage means. Alternatively the scheduler contains several separate, cyclical schedulers arranged in a queue, and each being dedicated to a single traffic service.

With this method and arrangement, the transmission efficiency across the shared link can be substantially improved because no exchange of queue information is necessary for the scheduled services. Furthermore, the delays and delay variations experienced by circuit emulated and other delay sensitive traffic can be kept at a minimum. By using multiple cyclical storage means, one dedicated to each service, both the implementation of the scheduler and its expansion to accommodate additional delay sensitive services is rendered exceptionally simple.

BRIEF DESCRIPTION OF THE DRAWINGS

Further objects and advantages of the present invention will become apparent from the following description of the preferred embodiments that are given by way of example with reference to the accompanying drawings, in which:

Fig. 1 schematically depicts a point to multipoint link of a multiplexed packet switched network,

Fig. 2 shows a first embodiment of a scheduler according to the invention, and

Fig. 3 shows a further embodiment of a scheduler according to the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

A typical multipoint-to-point link in a TDMA packet switched communication network is illustrated in Fig. 1. Several multipoint nodes 20 communicate via a single, shared bi-directional channel with a central node 10. Only three nodes are shown in the figure, however it is to be understood that further nodes may be connected above and below the illustrated nodes. The channel 100 may be a fixed electrical or electro-optical link or an air interface.

As in many conventional media access control (MAC) protocols, the central node 10 determines which time slots on the shared channel 100 will be allocated to which services conveyed by the multipoint nodes 20. This is accomplished by the central node 10 issuing permits for a time slot in the future information stream.

For delay insensitive packet switched services, such as data transmission services, the central node 10 holds a record of the number of packets of information awaiting transmission at each of the multipoint nodes 20. This is represented in the figure by a queue record memory 11. The queue records are updated by means of permit requests, or equivalent queue size reports, sent by the multipoint nodes 20 to the central node 10. When the record of a particular multipoint node 20 indicates that packets are awaiting transmission, the central node 10 issues a corresponding permit defining which time slot the multipoint node 20 in question may utilise for transmitting the packet. When more packets are awaiting transmission than can be dealt with immediately, the central node 10 determines which service has priority. This prioritising information is likewise obtained from the multipoint nodes 20 on invitation from the central node 10.

In the present example, however, it is assumed that the traffic A, B, C conveyed by the three illustrated multipoint nodes 20 is delay sensitive, i.e. has strict requirements as to allowable delays and delay variations. This may include, but is not limited to, circuit emulated traffic, voice or interactive video traffic. Furthermore, each service A, B, C has different frame lengths.

In accordance with the present invention, these services are not handled in the same way as delay insensitive services. While the central node 10 issues permits to the services A, B, C, these are not issued in response to information obtained from the multipoint nodes 20, but are instead scheduled by the central node 10 at regular intervals. In other words, no exchange of information takes place between the nodes. The issuing of permits is unsolicited. This is performed in accordance with the present invention with the aid of a scheduler.

Fig. 2 shows a scheduler 30 according to a first embodiment of the invention. The scheduler 30 can be viewed as a circular memory divided into cells 31, each of which represents a time slot on the shared channel. It will be understood, however, that the scheduler 30 may be implemented using one or more modulo counters, memories or registers with pointers, or similar means which allow information to be obtained at predetermined regular intervals. The scheduler 30 is either incorporated in the central node 10 or is arranged such that a controller of the central node 10 (not shown) that is adapted to handle the allocation of time slots across the shared link can easily access the scheduler 30.

Prior to operation, or upon the introduction of a new delay sensitive service, the scheduler 30 is programmed. This may be done by the central node 10 using its controller, or may be performed using external means. The scheduler 30 in Fig. 2 is already programmed with markers in some of the cells 31 that serve to indicate which service is to be allocated a permit for which time slot in the future traffic stream. In the present example these markers are represented by the letters A, B and C for the three delay sensitive services of the same names, respectively. The scheduler 30 does not define the slot as an absolute time, but rather the slot interval. The different services, A, B and C, have different frame formats and lengths determined by the transmission bit rate and denoted in the figure by a, b and c, respectively. In the figure, the frame length a of the service A is given as the equivalent of 7 time slots on the shared channel; the frame length b of service B is 8 time slots, and the frame length c of the service C is 11 time slots. In order that the permit intervals for each of these services can be respected, the total length of the scheduler 30 is a multiple of all the frame lengths in terms of time slots. While the scheduler 30 illustrated in Fig. 2 is somewhat truncated, it is assumed to have a length equal to the smallest multiple of a, b and c, that is 616 cells 31. It will be understood

that the permit intervals between different services will also take account of possible path delay differences between the various multipoint nodes 20 and the central node 10.

5 In operation, the central node 10, or its controller, consults the cells 31 of scheduler 30 in succession to determine whether a permit is scheduled for one of the services A, B or C for each time slot. If the cell 31 consulted is marked, a permit will be issued to the corresponding multipoint node 20 carrying the selected service. If no permit is scheduled for a particular time slot, the central
10 node 10 may issue a permit to one of the other delay insensitive services which require access to the shared medium on the basis of queue size information stored at the node 10, as described above. Once all the cells 31 of the scheduler 30 have been consulted in turn, the pointer or its equivalent returns to the beginning again.

It will be apparent that when services with different frame formats and lengths are to be scheduled at predetermined constant intervals over the same link, there will be incidences when permits for two or more different services will be scheduled in the same time slot. In this case, the central node 10, or rather
20 its controller, will allocate the time slot to one of the services and cause the frames for the remaining services to be delayed by one or more time slots. The allocation of time slots for conflicting permit schedules may be decided on the basis of prioritising information relating to the services, whereby one service will consistently be handled before the others. Alternatively, the central node 10 may treat all services with equal priority by giving one service priority in the first incidence of conflict and another service priority in the next
25 incidence. Alternatively, the scheduler 30 could be programmed such that only one permit is allocated to each cell, however this would necessarily introduce a firm priority for one or other of the services.

5 The scheduler 30 described with reference to Fig. 2 can deal efficiently with the access of circuit emulated delay sensitive traffic over a shared link of a packet switched network when the number of services with different frame lengths is not too high. However, when several different services, many with different frame lengths are to be scheduled on a shared medium, the length of the scheduler will become prohibitively large. Furthermore the scheduler 30 of Fig. 2 does not lend itself easily to expansion. It is complex to incorporate additional services, since new permit schedules will effect both the total length and the content of the scheduler 30.

10 This problem is overcome by the scheduler 40 illustrated in Fig. 3 according to a second embodiment of the invention.

15 In place of the single circular memory of the first embodiment of the invention, this scheduler 40 includes several short schedule memories 41 arranged in a queue. Each schedule memory 41 contains information for services having the same frame format and length. In Fig. 3 the topmost schedule memory 41 contains scheduling information for the service A, the next schedule memory 41 in descending order contains information relating to service B, and the final schedule memory 41 contains permit scheduling information for service C. The maximum length required for this scheduler is thus $a+b+c$ cells 42. Furthermore, since several services could be scheduled using each single queue, provided they have the same frame length the total length of the scheduler will not change. The cells 42 of each schedule memory 20 41 represent coincident time slots. As indicated by the arrows in Fig. 3, the central node 10, or its controller, thus consults the cells 42 of each of the schedule memories 41 representing the current time slot in turn, and determines which, if any, of the services is to be sent a permit for the specified

time slot.

As for the scheduler structure 30 of the first embodiment, some conflict is inevitable between the scheduled permits. Accordingly, when permit allocation markers are present in more than one cell 42 representing the current time slot, the central node controller will decide which service will be handled preferentially. This is done on the basis of prioritising information obtained previously from the corresponding multipoint nodes 20.

Alternatively, when no service is defined as having priority, or all services stipulate similar delays and delay variations to be respected, priority will be allocated by the central node controller to each service in turn as each conflict arises. Also, as for the previous embodiment, the schedule memories 40 may be programmed with a specified priority hierarchy causing some services to be consistently delayed in favour of another so that decisions do not need to be made by the central node 10, or its controller.

The scheduling of the transmission of some delay sensitive services over a shared link in a packet switched network in the manner described above can lead to good efficiency and allow the delays and delay variations to be kept small. However, the delays and delay variations can be reduced still further if the scheduling of a service is arranged to correspond with the assembly of cells. Circuit emulated traffic transmitted over a packet switched network may use only part of a cell for traffic information so as to reduce the time needed for assembly of the cell. For example, to reduce the time needed to assemble an ATM cell for a 64kbps telephone service, the ATM cell may contain only 24 bytes of traffic information with the 24 remaining bytes being filled with dummy data. The delay experienced by this service can be substantially reduced, when the scheduling of packet transmission using the scheduler 30, 40 described above is synchronised to coincide with the assembly of the cell